

POLICY MANAGEMENT FOR DISTRIBUTED COMPUTING AND A METHOD FOR AGING STATISTICS

5 **FIELD OF INVENTION**

This invention relates to a system and method for managing policy decisions for a network of cooperating computers and to a method for aging statistics.

10 **BACKGROUND OF INVENTION**

Prior art methods that implement policy for computer systems that are homogenous in architecture assigned work on a platform specific basis or an operating system basis. For example, the IBM S/390 workload manager uses 390 system platform specific statistics, such as, multi-programming level, virtual storage, expanded storage, to make decisions on where to place work. Prior art methods of policy management have required advance knowledge of how much CPU time or memory an application needs to run to efficiently assign the application in a cluster of computers and take advantage of the cluster resources. Prior art methods of policy management also have created an affinity between certain types of work and a specific computer.

Statistics used by prior art policy management methods have been updated on a periodic basis, whether or not new values have been received. This updating has been scheduled by a timer that signals the update times. This causes additional path length, CPU cycles and concerns about recovery, such as, failure of the timer to signal.

With the advent of computer networks and the distribution of work among the computers, there is a need for a policy manager that can manage policy independent of the architecture of the computers connected in the network.

There is also a need for a method and system for aging statistics that are used in the policy management process.

5 **SUMMARY OF INVENTION**

The present invention satisfies the aforementioned needs with a policy management system and method that manages the availability of a plurality of cooperating computers connected in a network to do work. The method identifies a set of specific ones of the plurality of cooperating computers as
10 available resources for the performance of the work. Performance metrics are derived from performance related values of the plurality of cooperating computers. Based on the performance metrics, the set of specific ones of the cooperating computers are changed to thereby enable the allocation of work to cooperating computers that, from a performance standpoint, can do the work
15 in the shortest possible time.

The system and method of the invention can be used with cooperating computers that are either homogenous or heterogeneous in architecture or that employ the same or diverse operating systems. The performance related
20 metrics, for example, include response times or queue delays of a cooperating computer.

The system and method of policy management is flexible to allow as much policy management as desired to be delegated from a system policy
25 manager to a local or cluster manager. The method may request a manager of a cluster to accept additional work or to give up pending work based on the performance related metrics or to start more work or to run more pieces of an application on one or more of the cooperating computers of the cluster.

30 According to an aspect of the present invention, the performance metrics are derived or updated only when a new value has been received or a

request has been made to view the data. This saves CPU cycles that were used in prior methods that performed updates periodically, whether or not new values had been received since the last update.

5 According to another aspect of the invention, the performance related values are received over a series of time intervals. The performance metrics are derived for n periods, of which the performance metrics of the n th period thereof includes an aggregate of the performance metrics for the current interval plus $n-1$ of the preceding time intervals. The performance metrics of
10 the n th or last period of a preceding interval are discarded during a current interval.

 The performance metrics are formed as a data structure having n rows that contain the performance metrics of the n periods. The performance
15 metrics of the n th row of a preceding interval are discarded during a current interval. The n th row of the preceding interval is used as a first row in the current interval and the remaining rows are shifted down one row position.

 According to a further aspect of the present invention, each of the
20 performance metrics includes only a number and average of values received.

BRIEF DESCRIPTION OF THE DRAWINGS

 Other and further objects, advantages and features of the present invention will be understood by reference to the following specification in
25 conjunction with the accompanying drawings, in which like reference characters denote like elements of structure and:

 FIG. 1 is a block diagram of a system in which the policy management system and method the present invention is employed;

30

FIG. 2 is a block diagram of the managing computer of the FIG. 1 system;

FIG. 3 is a flow diagram of the policy manager program of the FIG. 2 managing computer;

FIG. 4 is a flow diagram of the performance values and metrics formation process of the FIG. 2 managing computer;

FIG. 5 depicts a data structure of the performance metrics formed by the process of FIG. 4;

FIG. 6 is a flow diagram of the receive new values portion of the flow diagram of FIG. 4; and

FIG. 7 is a flow diagram of the update table portion of the flow diagram of FIG. 4.

DESCRIPTION OF THE PREFERRED EMBODIMENT

Referring to FIG. 1, a distributed computing system 20 includes a plurality of cooperating network computers 22-1 through 22-N, a managing computer 24, an intranet 26, an internet 28 and a cluster 30. Intranet 26 interconnects cooperating network computers 22-1 through 22-N, cluster 30 and managing computer 24 in a distributing computing network. Intranet 26 is generally a network that is internal to an organization. The present invention contemplates that other computers outside the organization may be included in the distributed computing network. To this end, internet 28 serves to interconnect these other computers (not shown) with managing computer 24, cluster 30 and cooperating network computers 22-1 through 22-N. Other non-network computers (not shown) can also communicate with system 20 via internet 28.

Cluster 30 includes a managing computer 36 and two cooperating network computers 32 and 34. Cooperating network computers 32 and 34 and managing computer 36 are each connected with intranet 26. Cluster 30 may be considered a node in system 20. Although two cooperating network computers 32 and 34 are shown, it will be apparent to those skilled in the art that cluster 30 may include more or less cooperating computers. It will also be apparent to those skilled in the art that system 20 may include additional clusters.

10

Managing computer 24 manages the policy concerning allocation of work (applications) and distribution of that work among cooperating computers 22-1 through 22-N and cluster 30. Managing computer 36 manages the policy and distribution of work among cooperating network computers 32 and 34. It will be apparent to those skilled in the art that managing computers 24 and 26 are shown as single computers by way of example and that policy management and work distribution may be functionally distributed among two or more computers.

20

Referring to FIG. 2, managing computer 24 includes a central processing unit (CPU) 40, an input/output (I/O) units section 42, a communications unit 44 and a memory 46. Communications unit 44 serves to interconnect managing computer with intranet 26 for communication with cooperating computers 22-1 through 22-N and cluster 30. Memory 46 includes an operating system program 48, a distribution manager program 52, a policy manager program 54, a metrics program 56 and a metrics data structure 58.

25

Policy manager program 54 and metrics program 56 control CPU 40 to develop and update metrics data structure 58 with performance related metrics of cooperating network computers 32, 34 and 22-1 through 22-N.

30

Distribution manager 52 uses metrics data structure 58 to control CPU 40 to distribute work among cooperating network computers 32, 34 and 22-1 through 22-N. Dependent on the allocation of policy management responsibility among managing computer 24 and managing computer 36, managing computer 36 may also have access to metrics data structure 58. It will be apparent to those skilled in the art that although metrics program 56 and metrics data structure 58 are shown as separate modules, each could be incorporated into policy manager program 54. It will be apparent to those skilled in the art that managing computer 36 may have the same or similar structure and software as managing computer 24.

Software stored in memory 46, including policy manager program 54, metrics program 56 may be loaded or down loaded from a memory medium 60 to memory 46.

Referring to FIG. 3, policy manager program 54 at step 70 identifies a set of resources that are available for work. For example, the specific ones of cooperating network computers 32, 34 or 22-1 through 22-N that are available for work are identified. This identification is available for access by distribution manager program 52. Step 72 examines the state of system 20 including metrics data structure 58. Step 74 determines if there is any need to change the available resources based on the metrics data structure 58. If no change is needed, step 76 causes policy manager program 54 to wait and then step 72 is repeated. If step 74 determines there is a need for change, step 78 either increases or decreases the current set of available resources. After step 78 is completed, step 76 causes policy manager program 54 to wait and then step 72 is repeated.

Referring to FIG. 4, metrics program 56 starts with step 80 that initializes based on provided parameters, such as, the location of basic components. Step 82 then causes metrics program 56 to wait until a request

is received. If the request is for a report, the metrics contained in metric data structure 58 are updated by step 86. Step 88 then issues the requested report and control returns to step 82. If the request is to provide new values, step 90 derives the metrics from the new values. The derived metrics are then used by step 92 to update metric data structure 58. Control is then passed to step 82. Metrics program 56 performs the foregoing steps for each cooperating network computer that is connected in system 20, unless managing computer 36 is responsible to track the performance related values of cooperating network computers 32 and 34.

10

Referring to FIG. 5, metrics data structure 58 is shown in the form of a table having rows 100, 102, 104 and 106. Rows 100, 102, 104 and 106 each contain a performance related metric for a different value reporting interval. According to one aspect of the invention, the performance related value may be either the response time or queue delay of a cooperating network computer. In FIG. 5, the performance related value is shown as response time. The derived metrics in each row consist of a number of values received for that reporting interval and an average of the reported values. Thus, row 100 is shown with a total number of 7 values received thus far in a current reporting interval with an average response time of 160 milliseconds (ms). Thus, only two metrics are needed for each reporting interval. Row 102 has a reporting interval that consists of the previous interval plus the current interval. Row 104 has a reporting interval of the two previous intervals plus the current interval. Row 206 has a reporting interval of the three preceding intervals plus the current interval. It will be apparent to those skilled in the art that there can be more or less rows than the four rows 100, 102, 104 and 106.

Referring to FIG. 6, aggregate values step 90 of FIG. 4 is shown in detail. Step 110 waits for a new value. Step 112 receives a new value. Step 114 determines if the current value reporting interval has expired. If so, step 116 updates rows 100, 102, 104 and 106 of metrics data structure 58 and

passes control to step 110. If not, step 118 updates only the current row 100. By updating only the current row as each new value is received, computation time is conserved. Step 120 increments the values number and step 122 calculates a new average. For example, the new average is the old average plus the difference of the current value and the old average divided by the new number of values. After step 122, control is passed to step 110.

Referring to FIG. 7, update table rows step 116 of FIG. 6 is shown in detail. Step 130 rotates table rows 100, 102, 104 and 106 of metrics data structure 58 of FIG. 5. Step 132 discards the metrics contained in the nth row, which is row 106 in FIG. 5. Row 106 is then used as the first row for the metrics to be derived in the next current interval. Step 134 then shifts all remaining rows down one position. Thus, for the next current interval, row 106 will contain the metrics for the current interval. Row 100 will contain the metrics for the previous interval plus that current interval. Row 102 will contain the metrics for the two previous intervals plus the current interval. Row 104 will contain the metrics for the three previous intervals plus the current interval.

Policy manager program 54 runs independently of distribution and execution of work within system 20. That is, the process of policy management is disconnected from the active work processes that are being managed. This has the effect of reducing overhead of system 20, thereby allowing for maximum scalability. Policy manager program 54 contains no process that establishes an affinity between a specific work item or type of work and a specific cooperating network computer.

Policy manager program 54 oversees system 20 and the work execution process without interfering with the production cycles of system 20. This is accomplished by having policy manager program 54 run separately from distribution manager program 52 and other software that controls execution of work by system 20. This allows policy manager program to

monitor the execution of work and to extract decision making information, such as performance related values, while minimally impacting the ongoing work process.

5 Policy manager program 54 and managing computer 24 are a central policy manager that monitors system 20. In order to achieve fault tolerance as well as limit the load on any one policy manager, managing computer 24 can communicate with a plurality of local or cluster policy managers, such as, managing computer 36, throughout system 20. The task of each local
10 managing computer is to enforce local policy, while meeting the goals of the global system through communication with other local policy managers and the central managing computer 24. Each cluster policy manager keeps track of the current policy and state of its cluster. By managing policy local to a cluster, the bottleneck of funneling all local policy decisions through a single
15 central policy manager is avoided. Clusters, such as cluster 30, can be partitioned based on functionality, proximity or any arbitrary consideration.

The policy management system and method of the present invention allows for more than one method for handling the distribution of policy
20 management among managing computer 24 and cluster managing computers, such as, managing computer 36. In one aspect, managing computer 24 can view cluster 30 as a single node and have no knowledge of cooperating network computers 32 and 34. In another aspect of the invention, managing computer 24 must approve of all decisions made by managing computer 36
25 and, thus, has first hand knowledge of system 20. A combination of these two aspects allows the cluster managing computer 36 to make local decisions about its resource management, while central managing computer 24, as needed, has access to the state of cluster 30 and cooperating network computers 32 and 34.

30

The present invention having been thus described with particular reference to the preferred forms thereof, it will be obvious that various changes and modifications may be made therein without departing from the spirit and scope of the present invention as defined in the appended claims.